


*Machine Learning*  
*CS4824/ECE4424*  
*Summer 2022*

**INSTRUCTOR: HONGJIE CHEN**  
**MAY 23RD 2022**



# About Hongjie Chen

- Nicknamed Jeff, a doctoral student in CS
- 4 years research experience
- Time-series and Graphs
  - FUN
- Working towards my degree
- My homepage: <https://people.cs.vt.edu/~jeffchan/>



# Important Links

- Homepage: <https://people.cs.vt.edu/~jeffchan/teaching/CS4824/index.html>
- Piazza: <https://piazza.com/class/l2gnhof3rrz3ax>
- Class zoom: <https://virginiatech.zoom.us/j/89545716672>
- Office hour zoom: <https://virginiatech.zoom.us/j/2023032020>
- TA office hour zoom: <https://virginiatech.zoom.us/my/aahuja>
- Canvas:
  - <https://canvas.vt.edu/courses/151640> (CS 4824)
  - <https://canvas.vt.edu/courses/151551> (ECE 4424)

# *Welcome Again!*

- What is Machine Learning?
- Logistics
- How to succeed in this class?

# *Discussion*

- What are some machine learning applications?

# *Discussion*

- What are some machine learning applications?
- For example
  - The face recognition function of the camera on your phone
  - Friends/News/Products/Videos Recommendation
  - Hey Siri/Alexa

# Definition: Machine Learning

- A computer program: learn from experience  $E$  with respect to some class of tasks  $T$  and performance measure  $P$ , if its performance at tasks in  $T$ , as measured by  $P$ , improves with experience  $E$ . (*Machine Learning, Tom Mitchell, McGraw Hill, 1997*)
- For example...
  - Building a face recognition program

# *Logistic*

- **What this course covers:** fundamental topics in Machine Learning by introducing key problems, intuitions to solutions, mathematical foundations, and realistic applications.
- **Goal:** After this course, students should be able to understand, recognize, analyze, explain, implement, and apply the covered models for research or industrial use.



# *This class DOES NOT target*

- Teaching the most edge-cutting ML methods and models;
  - Advanced Machine Learning class does
  - 5000/6000 level
- Teaching specific softwares or libraries for ML implementation;
  - scikit-learn
  - Tensorflow, Keras, DL4J, CNTK, Pytorch, MXNet, ...

# Topics\*

\*Tentative and subject to change

- **Basics of Statistical Learning:** Loss functions, MLE, MAP, Bayesian estimation, bias-variance tradeoff, overfitting, regularization, cross-validation
- **Supervised Learning:** Decision Trees, Naïve Bayes, Logistic Regression, Kernels and Kernel Regression, Support Vector Machines, Neural Networks
- **Unsupervised Learning:** EM, Clustering (K-means, Gaussian Mixture)
- **Graphical Models:** Bayesian Networks, Hidden Markov Models
- **Deep Learning:** Convolutional Neural Networks, Recurrent Neural Networks, Attention and Transformer Networks, Autoencoders, Variational Autoencoders, Generative Adversarial Networks
- **Reinforcement Learning:** Markov Decision Process, Value Iteration, Policy Iteration, Q-Learning
- **Machine Learning Applications:** Time-series, Graph Machine Learning

# Textbooks

- Not required.
- Optional reference books (free available online)
  - [Machine Learning: a Probabilistic Perspective](#), Kevin Murphy, MIT Press, 2012
  - [Pattern Recognition and Machine Learning](#), Christopher Bishop, Springer, 2006
  - [The Elements of Statistical Learning](#), Trevor Hastie, Robert Tibshirani, and Jerome Friedman, Springer, 2009
  - [Deep Learning](#), Ian Goodfellow, Yoshua Bengio, and Aaron Courville, MIT Press, 2016
  - [Reinforcement Learning: An Introduction](#), Richard S. Sutton and Andrew G. Barto, MIT Press, 2018

# Prerequisites

- Ability to deal with abstract mathematical concepts.
- Probability and Statistics (Get a reference book)
  - Basic concepts of probability including **random variables**, expectation, **chain rule conditional distribution**, **Bayes rule**, **likelihood**, prior probability, densities, marginalization, moments, etc.
- Calculus and Linear Algebra (Get a reference book)
  - Matrix multiplication, multivariate derivatives, **chain rule**.
- Algorithms
  - Basic data structures, complexity analysis.
- Programming
  - Heavy on Python, but not hard :)

# Homework Assignments and Grading

- HW assignments
  - 5 individual coding assignments (**19% x 5 = 95%**).
  - Python, Jupyter Notebook, (Google Colab)
  - Late submission with no penalty, one ticket, once used it's gone
  - Late submission with penalty (50% off if submitting before next midnight), otherwise 0pt
  - For late submission, students must reach out to the instructor Hongjie Chen
  - HW 5 must be submitted on time
  - **\*Start early!** Due at ET 11:59pm ~~Last minute submission~~
- Class participation (**5%**)
  - Contribute to discussions on Piazza
  - Engage in class: Q & A
- Final letter grade
  - A: 93.3%–100%, A-: 90.0%–93.3%, B+:86.6%–90.0%, B:83.3%–86.6%
  - B-:80.0%–83.3%, C+:76.6%–80.0%, C:73.3%–76.6%, C-:70.0%–73.3%
  - D+:66.6%–70.0%, D:63.3%–66.6%, D-:60.0%–63.3%, F:00.0%–60.0%

# *Policies*

- Regrading Request
  - If you find a grading error, email TA within 3 days of the grade release day.
- Honor Code
  - All assignments are individual assignments
  - Zero-tolerance on plagiarism. Honor Code Council.
- Principles of Community
  - Because the course will include in-class discussions, we will adhere to Virginia Tech Principles of Community.
- Accessibility
  - If any student needs special accommodations because of any disabilities, please contact the instructor during the first week of classes.
  - Such students are encouraged to work with The Office of Services for Students with Disabilities to help coordinate accessibility arrangements.

# *Stay in touch*

- Office hours\*: Tuesday 10:00 am - noon at **Personal Zoom**
- TA office hours: Wed & Fri 3:00 pm - 4:30 pm
- Piazza: <https://piazza.com/class/l2gnhof3rrz3ax>
  - Must not post answer
- Email: [jeffchan@vt.edu](mailto:jeffchan@vt.edu)
- Phone Call?

# *Suggestions*

- Take notes, recommend writing them down
- Don't hesitate to ask for clarification, in class or after class
- Start doing homework early
- Preview and review
  
- Have questions? Ask.



# Todo

- Check the course webpage: <https://people.cs.vt.edu/~jeffchan/teaching/CS4824/index.html>
- Login to Piazza: <https://piazza.com/class/l2gnhof3rrz3ax>
- Let me know if you have any questions